

IN THE CLAIMS:

*Requested
7/28
8/30/07*

This listing of claims will replace all prior versions, and listings, of claims in the present application:

1. (Currently amended) A method for mirroring data on a first storage server and a second storage server, the method comprising:

queuing write commands in a temporary storage of the first storage server between consistency points, wherein each consistency point is an event characterized by executing the write commands queued in the temporary storage to write data corresponding to a file system of the first storage server to a local set of mass storage devices coupled to the first storage server;

in response to initiation of each consistency point, concurrently sending the write commands queued in the temporary storage from the first storage server to the local set of mass storage devices and to a remote set of mass storage devices coupled to the second storage server, wherein the memory blocks of the local and remote sets of mass storage devices include address blocks that each store a pointer to another block in use by the file system, and data blocks that store data for the file system;

updating memory blocks of the local and remote sets of mass storage devices based on the write commands; and

at an end of each consistency point, constructing a representation to reference each memory block of the local set of mass storage devices that is in use to represent the file system, the representation being in the form of a sub-tree structure that includes nodes that represent one

or more of the address blocks and one or more of the data blocks, the sub-tree structure being rooted in a root node that includes a pointer to the sub-tree structure; and

sending only said root node from the first storage server to the second storage server, said root node comprising information to allow reconstruction of the entire representation by the second storage server.

2-5. (Canceled)

6. (Currently amended) The method of claim 1, wherein updating the memory blocks is in accordance with a scheme in which a memory block referenced within a representation is not overwritten.

7. (Original) The method of claim 6, further comprising allowing read-only access to the remote set of mass storage devices while updating the memory blocks of the remote mass storage server.

8. (Currently amended) A method for operating a first storage server to mirror a file system of the first storage server on a second storage server, the method comprising:

receiving high-level write requests from a client served by the first storage server, the high-level write requests requiring changes to the file system;

generating block-level write commands to update memory blocks of a local set of mass storage devices coupled to the first storage server, based on the high-level write requests;

Request
DD
8/30/07

queuing the block-level write commands in a temporary storage of the first storage server between consistency points, wherein each consistency point is an event characterized by executing the write commands queued in the temporary storage to write data corresponding to the file system to the local set of mass storage devices; and

in response to initiation of each consistency point, concurrently sending the block-level write commands queued in the temporary storage from the first storage server to the local set of mass storage devices and to a remote set of mass storage devices coupled to the second storage server to cause the memory blocks of the local and remote sets of mass storage devices to be updated, wherein the memory blocks of the local and remote sets of mass storage devices include address blocks that each store a pointer to another block in use by the file system, and data blocks that store data for the file system; constructing a representation to reference each memory block of the local set of mass storage devices that is in use to represent the file system after the memory blocks have been updated, the representation being in the form of a sub-tree structure that includes nodes that represent one or more of the address blocks and one or more of the data blocks, the sub-tree structure being rooted in a root node that includes a pointer to the sub-tree structure; and sending only the root node from the first storage server to the second storage server, said root node comprising information to allow reconstruction of the entire representation by the second storage server.

9-12. (Canceled)

13-18. (Canceled)

19. (Currently amended) A storage server, comprising:

a processor; and

a memory coupled to the processor, the memory storing instructions which when executed by the processor, cause the storage server to perform a method comprising:

receiving high-level write requests from a client served by the storage server, the high-level write requests requiring changes to a file system of the storage server;

generating block-level write commands to update memory blocks of a local set of mass storage devices coupled to the storage server based on the high-level write requests;

queuing the block-level write commands in a temporary storage of the storage server between consistency points, wherein each consistency point is an event characterized by executing the write commands queued in the temporary storage to write data corresponding to the file system to the local set of mass storage devices;

in response to initiation of each consistency point, concurrently sending the block-level write commands queued in the temporary storage from the storage server to the local set of mass storage devices and to a remote set of mass storage devices coupled to another storage server to cause the memory blocks of the local and remote sets of mass storage devices to be updated, wherein the memory blocks of the local and remote sets of mass storage devices include address blocks that each store a pointer to another block in use by the file system, and data blocks that store data for the file system; constructing a representation to reference each memory block of the local set of mass storage devices that is in use to represent the file system, after the memory blocks have been updated, the representation being in the form of a sub-tree structure that includes nodes that represent one or more of the address blocks and one or more of the data

blocks, the sub-tree structure being rooted in a root node that includes a pointer to the sub-tree structure; and sending only said root node from said storage server to the other storage server, said root node comprising information to allow reconstruction of the entire representation by the other storage server.

20-23. (Canceled)

24-29. (Canceled)

30. (Currently amended) A computer-readable medium, having stored thereon a sequence of instructions, which when executed by a computer, cause the computer to perform a method comprising:

receiving high-level write requests from a client served by a first storage server, the high-level write requests requiring changes to a file system of the first storage server;

generating block-level write commands to update memory blocks of a local set of mass storage devices coupled to the first storage server, based on the high-level write requests;

queuing the block-level write commands in a temporary storage of the first storage server between consistency points, wherein each consistency point is an event characterized by executing the write commands queued in the temporary storage to write data corresponding to the file system to the local set of mass storage devices; and

Revised to 8/30/07

in response to initiation of each consistency point, concurrently sending the block-level write commands queued in the temporary storage from the first storage server to the local set of mass storage devices and to a remote set of mass storage devices coupled to a second storage server to cause the memory blocks of the local and remote sets of mass storage devices to be updated, wherein the memory blocks of the local and remote sets of mass storage devices include address blocks that each store a pointer to another block in use by the file system, and data blocks that store data for the file system; constructing a representation to reference each memory block of the local set of mass storage devices that is in use to represent the file system, after the memory blocks have been updated, the representation being in the form of a sub-tree structure that includes nodes that represent one or more of the address blocks and one or more of the data blocks, the sub-tree structure being rooted in a root node that includes a pointer to the sub-tree structure; and sending only said root node from the first storage server to the second storage server, said root node comprising information to allow reconstruction of the entire representation by the second storage server.

31. (Canceled)

32-33. (Canceled)

34-35. (Canceled)

36. (Currently amended) A method comprising:

*Requested
SS
8/30/07*

at each of a series of consistency points, constructing a representation of a file system, the representation being in the form of a sub-tree structure that includes a plurality of nodes and includes pointers to each memory block used to store a file system maintained by a first server, the sub-tree structure being rooted at a root node, said each memory block belonging to a first set of mass storage devices coupled locally to the first server, wherein each consistency point is an event characterized by, in response to initiation of the consistency point, sending block-level write commands, queued in a temporary storage of the first server between consistency points, concurrently to the first set of mass storage devices to update memory blocks of the first set of mass storage devices and to a second set of mass storage devices coupled locally to a second server to update memory blocks of the second set of mass storage devices; and

sending only said root node from the first server to the second server at each said consistency point, said root node comprising information to allow reconstruction of the entire representation by the second server.

37. (Canceled)

38. (Canceled)